Sports betting strategies: an experimental review

Matej Uhrín[1], Ondřej Hubáček[1], Gustav Šourek[1], and Filip Železný[1]

Czech Technical University, Prague

Abstract

We investigate the problem of optimal wealth allocation over predictive sports market's opportunities. We analyze the problem across diverse settings, utility targets, and the notion of optimality itself. We review existing literature to identify the most prominent approaches coming from the diverse sport and economic views on the problem, and provide some practical perspectives. Namely, we focus on the provably optimal geometric mean policy, typically referred to as the Kelly criterion, and Modern Portfolio Theory based approaches leveraging utility theory. From the joint perspective of decision theory, we discuss their unique properties, assumptions and, importantly, investigate effective heuristics and practical techniques to tackle their key common challenges, particularly the problem of uncertainty in the outcome probability estimates. Finally, we verify our findings on a large dataset of soccer records.

1 Introduction

In this work, we understand the game of betting as an investment decision making over probabilistic outcomes. We generally assume two sides acting in the game setting – a bookmaker (B) and a player model (M for "model"). Both the bookmaker and the player have their beliefs about the true probability distribution (P_R) over the outcomes. Prior to the outcome realization, the bookmaker sets up possible payoffs, representing the returns of investment for each possible outcome realization.

Example 1. *As a simplistic, artficial example of such a setting, imagine a game of fair coin tossing [7], where the payoffs are set up, for instance, as follows. If the coin ends up heads, our wealth grows by 50%, and if it is tails, our wealth shrinks by 40%.*

Generaly, the goal of the player is to maximize the expected profits while minimizing expected losses (see Section 2 for specific target utility definitions), assuming repeated investment over a prolonged period of time.

Player's betting system typically consists of two major components. A model, estimating the probability of each outcome, and a betting strategy, combining the probabilities given by the predictive model with the bookmaker's payoffs to determine how to split the stakes. In this work, we focus solely on the second component – the betting strategy.

1.1 Betting Opportunities

The opportunities in the game represent simply the set of probabilistic outcomes currently available for betting. We can generally cover each betting game setting via definition of a "payoff matrix" $\boldsymbol{R}$, the columns of which represent different opportunities available, and rows represent different outcome realizations. Each element in $\boldsymbol{R}$ then represents a single payoff from each opportunity realization. Additionally, we include a risk-free "cash" opportunity $\boldsymbol{c}$, which allows our strategy to put money "aside". Also, $\boldsymbol{c}$ allows to model situations where leaving money aside can cost small fraction of $\boldsymbol{c}$ in every turn ("inflation"), or possibility to increase with some small interest rate (e.g. in a bank account).

The betting strategy will thus allocate the current wealth among n opportunities, $n-1$ of which are the risky, probabilistic assets, with 1 opportunity being the added risk-free cash asset. Let us further assume that there are K possible worlds, i.e. K possible joint outcome realizations, in our probabilistic game. Each opportunity can then be fully specified using an asset vector $\boldsymbol{a_i}$ with returns $r_{i,k}$ for each of the K probabilistic outcomes $p_1, p_2, ..., p_K$ [1].

$$\boldsymbol{R} = \begin{bmatrix} \boldsymbol{a_1} & \boldsymbol{a_2} & ... & \boldsymbol{a_{n-1}} & \boldsymbol{c} \end{bmatrix}, \ where \ \boldsymbol{a_i} = \begin{bmatrix} r_{i,1} \\ r_{i,2} \\ ... \\ r_{i,K} \end{bmatrix}, \ \boldsymbol{c} = \begin{bmatrix} 1 \\ 1 \\ ... \\ 1 \end{bmatrix} \tag{1}$$

1.2 Betting strategy

Betting strategy takes bookmaker's odds together with model's estimates, and outputs portions of wealth, i.e. the bets, to be waged on each of the betting opportunities. We can formally define the betting strategy f for n betting opportunities as follows:

$$f : P_M^n \times O^n \rightarrow B^n \tag{2}$$

where P_M is the probability distribution over the outcome realizations given by model, O are the bookmaker's odds, and B are the portions of wealth staked. We will further refer to the vector of wagers $\boldsymbol{b}$ as "portfolio":

$$\boldsymbol{b} = \begin{bmatrix} b_1, b_2, ..., b_n \end{bmatrix}^T \tag{3}$$

where b_i stands for wealth portion allocated to i-th opportunity.

1.3 Betting dynamic

A betting dynamic is part of the game setting for the presented opportunities. The dynamic specifies the investment behavior as the opportunities are being repeated in time, determining the progression of wealth. For that goal, we will use the following definitions [7]:

- $W(t)$ is wealth in time t.
- δt is a regular time interval, e.g. 1 week.
- $r(t)$ is a return/payoff function in time t.

For example, the return function for the coin toss game from Example 1 can be defined as follows [7]:

$$r(t) = \begin{cases} 1.5 & \text{with probability } 1/2 \\ 0.6 & \text{with probability } 1/2 \end{cases} \tag{4}$$

We further assume that if we accept, we are obliged to play this game for a longer period of time (e.g. every week for 2 years). Generally, the investment behavior in this game can be either additive or multiplicative, determining the dynamic of our wealth progression.

1.3.1 Additive dynamic

Additive dynamic is simply a unit based investment, meaning that we invest a single unit (e.g. 1 dollar) at every single time step δt. Our wealth progression in time t is hence defined as:

$$W(t) = r(t) + W(t - \delta t) \tag{5}$$

Taking the above mentioned coin toss example, the "growth rate" of player's wealth will approximately be the expected value $ev = \frac{1}{2} \cdot 1.5 + \frac{1}{2} \cdot 0.6 = 1.05$, in other words wealth growth rate under additive dynamic is ergodic [7].

1.3.2 Multiplicative dynamic

In this case, we are continuously reinvesting our current wealth over the set of the presented opportunities. Hence our wealth progression in time t is defined as follows:

$$W(t) = r(t) \cdot W(t - \delta t) \tag{6}$$

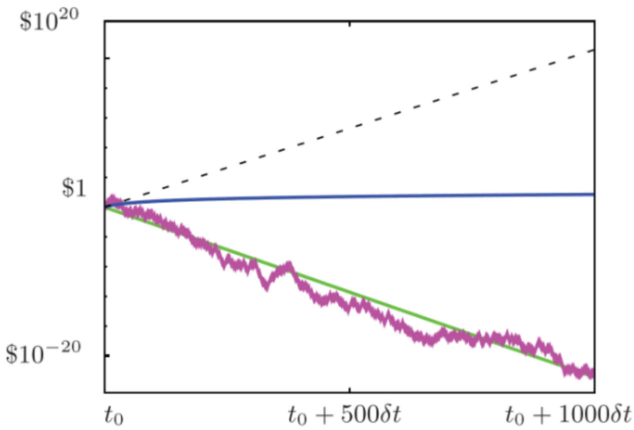

Figure 1: Multiplicative dynamic: Coin toss game. The magenta line represents median wealth trajectory in 1000 time steps of a coin toss game. The dashed line is the expected value [7].

In the Figure 1 we can see that, under multiplicative dynamic, the expected value does not indicate what happens to a single player in the "long run" [7]. Instead, the relevant quantity to look for in this case is the geometric mean $\bar{r} = (1.5 \cdot 0.6)^{\frac{1}{2}} \approx 0.95$. In other words, wealth growth rate under multiplicative dynamic is non-ergodic [7].

1.3.3 Ergodic property

Ergodic property of a dynamic process is often dubbed as the "equality of averages", which stands for the time average [7] being equal to the ensemble average, i.e. the expected value. The growth rate of wealth is thus ergodic if its expected value is constant in time and its time average converges to this value with probability one [7].

2 Betting Strategies

In existing literature, the betting strategies range from informal "ad-hoc" approaches, such as betting according to the probability estimates, to the formal ones, represented mainly by Modern portfolio theory [6] and Kelly criterion [4].

2.1 Informal approaches

Here we list some of the informal betting strategies encountered in the literature for completeness of the review.

- Bet fixed amount on favorable odds (unif).
- Bet amount equal to the absolute discrepancy between probabilities predicted by the model and the bookmaker (abs disc bet).
- Bet amount equal to the relative discrepancy between probabilities predicted by the model and the bookmaker (rel disc bet).
- Bet amount equal to the estimated probability of winning (conf bet).
- Bet everything only on the maximum expected value opportunity (max ev).

While these informal betting methods are generally inferior to the formal approaches [2] and lack the necessary theoretical background, we do not investigate them experimentally much further in this work.

2.2 Modern Portfolio Theory

The idea behind Modern Portfolio Theory (MPT) is that portfolio $\boldsymbol{b_1}$ is superior to $\boldsymbol{b_2}$ if the expected gain $\mathbb{E}[g(\boldsymbol{b})]$ is at least as great.

$$\mathbb{E}[g(\boldsymbol{b_1})] \geq \mathbb{E}[g(\boldsymbol{b_2})] \tag{7}$$

and the risk, here generaly denoted through a risk measure r, is no greater [6].

$$r(\boldsymbol{b_1}) \leq r(\boldsymbol{b_2}) \tag{8}$$

This creates a partial ordering on the set of all available portfolios. Taking the portfolios that no other portfolio is superior to gives us the set of efficient portfolios Θ. In simple terms we maximize the following:

$$\mathbb{E}[gain] - \gamma \cdot risk \tag{9}$$

In the most common setup, the risk of a portfolio is measured by its variance defined through a covariance matrix $\boldsymbol{\Sigma}$. MPT can then be expressed as the following maximization problem:

$$\begin{aligned} \underset{\boldsymbol{b}}{\text{maximize}} \quad & \boldsymbol{\mu}^T \boldsymbol{b} - \gamma \boldsymbol{b}^T \Sigma \boldsymbol{b} \\ \text{subject to} \quad & \sum_{i=1}^{K} b_i = 1.0, \; b_i \geq 0 \end{aligned}$$

where $\boldsymbol{b}$ is portfolio, γ is risk aversion parameter and $\boldsymbol{\mu}$ is the expected gains vector of the offered opportunities.

2.2.1 Maximum Sharpe strategy

Sharpe ratio can be used as a criterion to choose from the set of efficient portfolios Θ. Sharpe ratio of a portfolio is defined as [2]:

$$\frac{r_p - r_f}{\sigma_p}$$

Where σ_p is standard deviation of portfolio return, r_p is expected return of the portfolio and r_f is a risk-free rate (we can neglect the risk free rate if there is no risk-free method how to appreciate the money). We can hence define a separate "MaxSharpe" strategy as follows:

$$\begin{aligned} \underset{\boldsymbol{b}}{\text{maximize}} \quad & \frac{\boldsymbol{\mu b}}{\sqrt{\boldsymbol{b}^T \boldsymbol{\Sigma b}}} \\ \text{subject to} \quad & \sum_{i=1}^{K} b_i = 1.0 \\ & b_i \geq 0 \end{aligned}$$

2.2.2 Criticism

The MPT approach is often criticised for two main reasons, [7].

1. The growth rate of wealth under multiplicative dynamic is non-ergodic. In other words, the expected return does not tell us what will really happen in the long run, hence maximizing it does not produce a truly long term-optimal reinvestment strategy, which we showcased in the coin toss example 1.

2. The definition of risk as variance is often disputed. In many domains, risk is not easy to define.

2.3 Growth optimal strategy

A different approach to betting is to first make the growth rate of wealth ergodic, under specified dynamic using appropriate transformation, and then take the expected value. This approach is famously known as the "Kelly Criterion", "Geometric Mean Policy" and under many other names.

For the multiplicative dynamic, the correct ergodicity transformation is the logarithm: $v(x) = \log(x)$ [7]. The growth optimal portfolio for a general case can hence be found by solving the following optimization problem:

$$\begin{aligned} \underset{\boldsymbol{b}}{\text{maximize}} \quad & \mathbb{E}[\log(\boldsymbol{R} \cdot \boldsymbol{b})] \\ \text{subject to} \quad & \sum_{i=1}^{K} b_i = 1.0, \; b_i \geq 0 \end{aligned}$$

The calculated portfolio is then growth-optimal under the following assumptions [4, 8, 7]:

1. True probability is known to the player.

2. There is a "long run" of "approximately similar" games.

3 Experiments

We showcase the above mentioned strategies on a dataset of 30000 soccer games from different leagues all over the world. The dataset consists of opening odds, closing odds, estimated outcome probabilities and results indicating the outcome realization in each game. The probabilities are modelled separately using a neural network model. Naturally, the true probabilities of the outcomes are unknown to both to the player and the bookmaker, which has a significant impact on how the strategies are applied, especially in the multiplicative dynamic scenario.

We have three outcomes per game and we always assume multiple simultaneous games happening simultaneously in each time step, which we will refer to as "round". We investigate results from both the additive and multiplicative perspective.

In [6] the author proposes multiple measures of dispersion (risk measures) that can possibly be used to evaluate the investment strategy such as variance, standard deviation and coefficient of variation. In our evaluation framework we choose the standard deviation of portfolio return as the risk measure.

3.1 Additive dynamic

If our goal would be solely to maximize the expected profit, then the solution would be a trivial Bayesian strategy – bet the whole wealth on the opportunity with highest expected value of profit from presented opportunities. However, even under additive dynamic, we want some level of guarantee of not loosing too much money, hence our criteria for selecting the betting strategy are the expected profit and also the risk.

To demonstrate the differences between the betting strategies we randomly sampled data representing six illustrative betting opportunities with positive expected value [3]. Then we applied the betting strategies and analyzed the differences in the wealth allocations by the different strategies. Results are summarized in Table 1 (we note that the growth optimal approach only makes sense under the multiplicative dynamic, hence we do not analyse it in this sub-section). The strategies based on the absolute and relative discrepancies between the probability estimates always prefer higher expected profit of the opportunity regardless of the variance of the profit. On the other hand, the strategy based on the confidence of the probability estimate always prefers lower risk, regardless of the return. The MaxSharpe strategy is looking for a compromise between these two approaches.

#	p_M	p_B	σ	ev	MaxEV	Unif	Abs_disc	Rel_disc	Conf	MaxSharpe
1	0.30	0.26	1.80	1.19	1.0	0.17	0.20	0.35	0.08	0.09
2	0.59	0.52	0.94	1.12	0	0.17	0.27	0.24	0.16	0.23
3	0.75	0.70	0.62	1.07	0	0.17	0.21	0.15	0.21	0.30
4	0.60	0.57	0.86	1.06	0	0.17	0.13	0.12	0.17	0.12
5	0.74	0.71	0.62	1.04	0	0.17	0.11	0.08	0.20	0.17
6	0.64	0.62	0.77	1.03	0	0.17	0.07	0.06	0.18	0.08

Table 1: Comparison of betting strategies on simulated betting opportunities. The columns p_M and p_B represent the probability estimates of the model and the bookmaker, respectively, σ and ev refer to standard deviation and expected value of the opportunity [3].

The paper [3] provides a simulation of of all the above mentioned informal strategies on a dataset of 5000 basketball games. The author comes to the conclusion that the only relevant strategies are MaxEV and MaxSharpe. While MaxEV directly maximizes the expected profit, it completely ignores the risk and the possibility of ruin. Hence, we do not consider MaxEV to be a "reasonable" betting strategy, and we will further focus solely on the MaxSharpe strategy in our analysis of additive dynamic.

3.1.1 Simultaneous games

In the real world, there is almost always multiple soccer games occuring in any given moment. Hence it is reasonable to assume and conduct our simulation in such a way so as to test our strategies under this assumption. The number of simultaneous games i.e. "round size" clearly affects the number of presented opportunities. The more soccer games we can bet on, the more opportunities are presented to us. In this subsection we investigate whether round size significantly affects the overall cumulative profit.

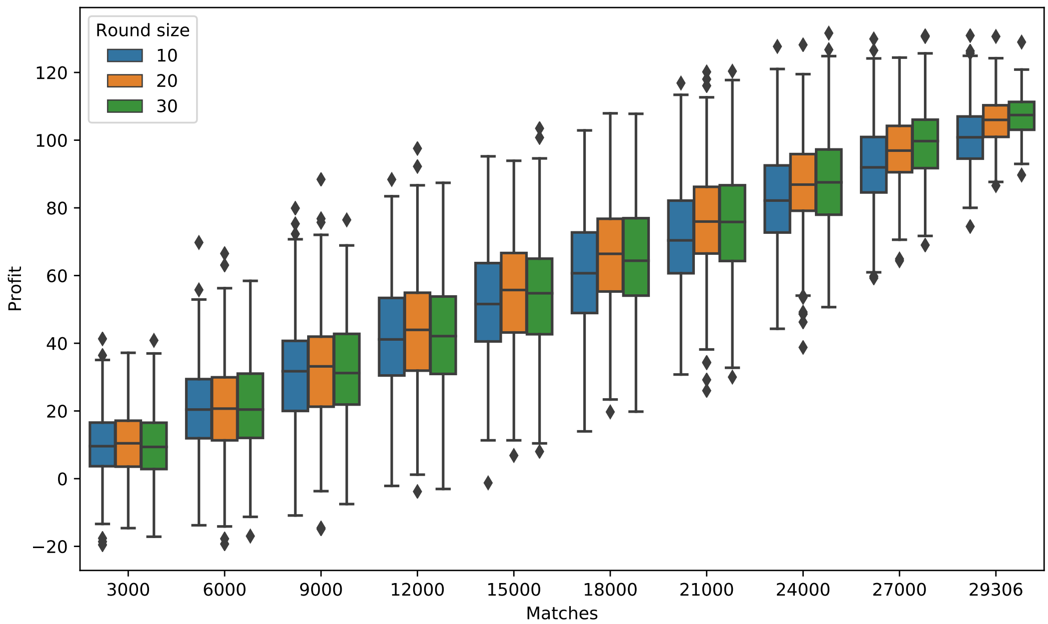

Figure 2: Simultaneous games: Effect of round size on profit under additive dynamic.

3.2 Multiplicative dynamic

In this section we focus on the reinvestment scenario. Our goal is hence to not only evaluate the presented oppoprtunities, but to evaluate them in context of the resources available to us, i.e. our "bank". We present two strategies that produced the best results in our experiments namely MaxSharpe strategy and the growth optimal Kelly criterion. We omit the informal approaches from the subsection 3.1 as they are insufficient for reinvestment scenario.

3.2.1 Fractional MaxSharpe

The main idea behind any "fractional" approach is to bet only a fraction ω of the calculated portfolio and leave the rest $1-\omega$ in the cash asset for security. Here we investigate results of the previously mentioned MaxSharpe strategy. We define a trade-off index ω for a portfolio as:

$$\boldsymbol{b_\omega} = \omega \boldsymbol{b_s} + (1-\omega)\boldsymbol{b_c} \tag{10}$$

where the $\boldsymbol{b_s}$ stands for portfolio suggested by MaxSharpe strategy and $\boldsymbol{b_c}$ is a portfolio where the only investment is a risk-free, "cash" investment.

Fraction, or a trade-off index, ω is now a lever between "growth" and "security" [5]. We hence fit the parameter ω on a training set and verify it on the testing set. We are looking for a reinvestment strategy that satisfies the following maximization problem across all the training dataset wealth trajectories.

$$\begin{aligned} \text{maximize} \quad & median(\boldsymbol{W_F}) \\ \text{subject to} \quad & Q_5 > 0.95 \end{aligned}$$

We are hence looking for a strategy that reaches the maximum median final wealth across all wealth trajectories, with the 5th percentile being the value below which 5% of all the wealth positions may be found. For the fractional MaxSharpe strategy this criterion yields a fraction $\omega = 0.11$

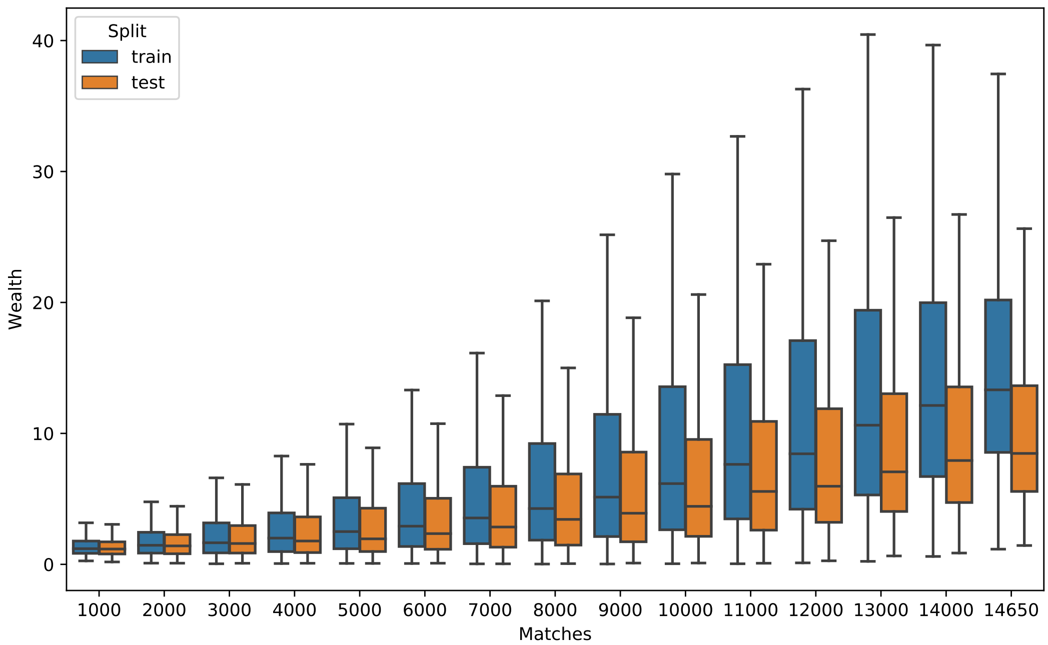

Figure 3: Results of the risk constrained growth optimal strategy on football dataset.

3.2.2 Risk constrained Kelly

In this case we adjust the growth optimal portfolio $\boldsymbol{b}$ using the following maximum drawdown constraint.

$$P(W^{MIN} < \alpha) \leq \beta \tag{11}$$

representing the probability of our wealth falling below α being at most β. This constraint is approximately satisfied if the following is satisfied [1]:

$$\mathbb{E}[(\boldsymbol{R} \cdot \boldsymbol{b})^{-\lambda}] \leq 1 \;\; where \;\; \lambda = \log(\beta)/\log(\alpha) \tag{12}$$

We hence fit the parameter λ on a training set and verify it on the test set using the same criterion as in the previous section. For the risk constrained Kelly the parameter is $\lambda = 9.4$.

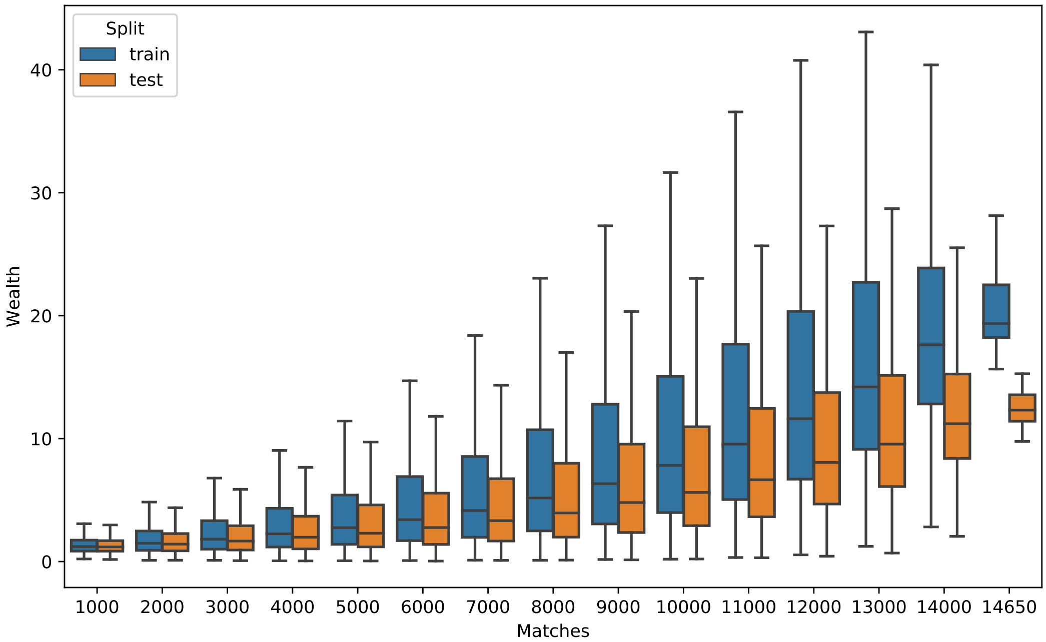

Figure 4: Results of the risk constrained growth optimal strategy on football dataset.

4 Results

We conducted experiments on a dataset of football games and our findings are as follows.

- The round size indeed does matter. Suppose that there are too few games taking place in a given day (e.g. 5). In such a case, the round is not representative enough and the strategy has too few opportunities to compare against each other, which leads to a lower diversificaiton of risk. On the other hand, the more games we are able to aggregate, the smaller the differences between the rounds, which as a result leads to a higher stability of the betting strategy.

- Both fractional MaxSharpe and risk constrained Kelly reach very similar results over the football dataset.

5 Conclusion

In this paper we review the most widely used betting strategies both theoretically and practically. From theoretical point of view we discuss the unique properties of the additive dynamic and the multiplicative dynamic, in other words the unit based and reinvestment scenario. Generally, we argue for the formal approaches as opposed to the informal ones often encountered in literature.

From the practical point of view, we showcase application of the approaches in a relevant sports domain, with several ideas on how to tackle one of its most difficult challenges – the uncertainty of probability estimates.

References

[1] Enzo Busseti, Ernest K Ryu, and Stephen Boyd. Risk-constrained kelly gambling. *arXiv preprint arXiv:1603.06183*, 2016.

[2] Ondrej Hubacek. Exploiting betting market inefficiencies with machine learning. Master's thesis, Czech Technical University, Faculty of Electrical Engineering, 2017.

[3] Ondřej Hubáček, Gustav Šourek, and Filip Železnỳ. Exploiting sports-betting market using machine learning. *International Journal of Forecasting*, 35(2):783–796, 2019.

[4] John L Kelly Jr. A new interpretation of information rate. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 25–34. World Scientific, 2011.

[5] LC MacLean, William T Ziemba, and George Blazenko. Growth versus security in dynamic investment analysis. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 331–354. World Scientific, 2011.

[6] Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

[7] Ole Peters and Murray Gell-Mann. Evaluating gambles using dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(2):023103, 2016.

[8] Edward O Thorp. The kelly criterion in blackjack sports betting, and the stock market. In *Handbook of asset and liability management*, pages 385–428. Elsevier, 2008.

ISBN 978-618-5036-53-9